arXiv:1407.0419v3 [math.OC] 4 Mar 2015

Conservative Signal Processing Architectures For Asynchronous, Distributed Optimization Part II: Example Systems

Thomas A. Baran and Tarek A. Lahlou
Digital Signal Processing Group
Massachusetts Institute of Technology

***Abstract*—This paper provides examples of various synchronous and asynchronous signal processing systems for performing optimization, utilizing the framework and elements developed in a preceding paper. The general strategy in that paper was to perform a linear transformation of stationarity conditions applicable to a class of convex and nonconvex optimization problems, resulting in algorithms that operate on a linear superposition of the associated primal and dual decision variables. The examples in this paper address various specific optimization problems including the LASSO problem, minimax-optimal filter design, the decentralized training of a support vector machine classifier, and sparse filter design for acoustic equalization. Where appropriate, multiple algorithms for solving the same optimization problem are presented, illustrating the use of the underlying framework in designing a variety of distinct classes of algorithms. The examples are accompanied by numerical simulation and a discussion of convergence.**

***Index Terms*—Asynchronous optimization, distributed optimization, conservation**

I. Introduction

This paper presents various classes of asynchronous, distributed optimization systems, demonstrating the use of the framework discussed in Part I [1]. The design and use of each class of systems is based upon the following strategy:

1) Write a reduced-form optimization problem, defined in [1].
2) Connect appropriate constitutive relations to interconnection elements, e.g. from Figs. 2-3 in [1], implementing the associated transformed stationarity conditions. Delay-free loops will generally result.
3) Break delay-free loops:
   a) For any constitutive relation that is a source element, perform algebraic simplification thereby incorporating the solution of the algebraic loop into the interconnection.
   b) Insert synchronous or asynchronous delays between the remaining constitutive relations and the interconnection.
4) Run the distributed system until it reaches a fixed point. The discussion in Section III, in conjunction with the system properties in Fig. 3 in [1], provide guidance in determining when convergence is ensured.
5) Read out the primal and dual decision variables a_i and b_i by multiplying the variables c_i and d_i by the inverses of the (2×2) matrices used in transforming the stationarity conditions.

II. Example Systems

Figs. 2-7 depict various asynchronous, distributed optimization algorithms implemented using the presented framework, specifically making use of the elements in Figs. 2-3 of Part I [1]. Figs. 2 and 3 in this paper illustrate two alternative implementations of systems for solving the LASSO problem. Figs. 4 and 5 depict two alternative implementations of systems for performing minimax-optimal FIR filter design. Fig. 6 depicts a support vector machine classifier trained using a decentralized algorithm generated using the presented framework. Fig. 7 illustrates an example of a nonconvex optimization algorithm aimed at the problem discussed in [2], in particular that of designing a sparse FIR filter for acoustic equalization. In Figs. 2-7, the asynchronous delay elements were numerically simulated using discrete-time sample-and-hold systems triggered by independent Bernoulli processes, with the probability of sampling being 0.1.

The authors wish to thank Analog Devices, Bose Corporation, and Texas Instruments for their support of innovative research at MIT and within the Digital Signal Processing Group.

III. Discussion of convergence

Fig. 1(a) summarizes the overall interconnection of elements composing the presented class of systems discussed in Part I [1], with those maps $m_k(\cdot)$ corresponding to source relationships being written separately. Figs. 1(b)-(d) illustrate a set of manipulations useful in analyzing convergence, with Fig. 1(b) specifically depicting a solution to the transformed stationarity conditions. The approach is to begin with the system in Fig. 1(a) and perform the additions and subtractions of $c_i^\star$ and $d_i^\star$ indicated in Fig. 1(c), obtaining Fig. 1(d) by identifying that Fig. 1(c) is a superposition of Figs. 1(b) and (d).

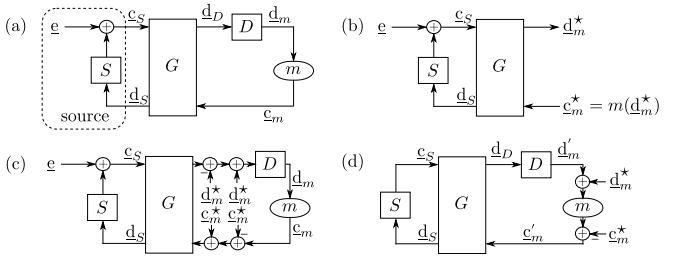

Fig. 1. (a) General description of the interconnection of elements used in the presented systems. (b)-(d) Manipulations performed in analyzing convergence.

There are various ways that the system in Fig. 1(d) can be used in determining sufficient conditions for convergence, a subset of which we outline here. Generally, arguments for convergence utilizing Fig. 1(d) involve identifying conditions for which $||\underline{\mathrm{d}}_D||$ in this figure is strictly less than $||\underline{\mathrm{d}}'_m||$, except at 0. Using the definition of a source element in [1] and the fact that G is a neutral map, i.e. an orthonormal matrix, we conclude from Fig. 1(d) that

$$||\underline{\mathrm{d}}_D|| \leq ||\underline{\mathrm{c}}'_m||. \tag{1}$$

If, for example, the solution to the transformed stationarity conditions $c_i^\star$ and $d_i^\star$ is known to be unique, and additionally if the collection of constitutive relations denoted $m(\cdot)$ is known to be dissipative about $\underline{\mathrm{d}}_m^\star$, then from Fig. 1(d) we conclude that $||\underline{\mathrm{c}}'_m|| < ||\underline{\mathrm{d}}'_m||$ except at 0, resulting in

$$||\underline{\mathrm{d}}_D|| < ||\underline{\mathrm{d}}'_m|| \tag{2}$$

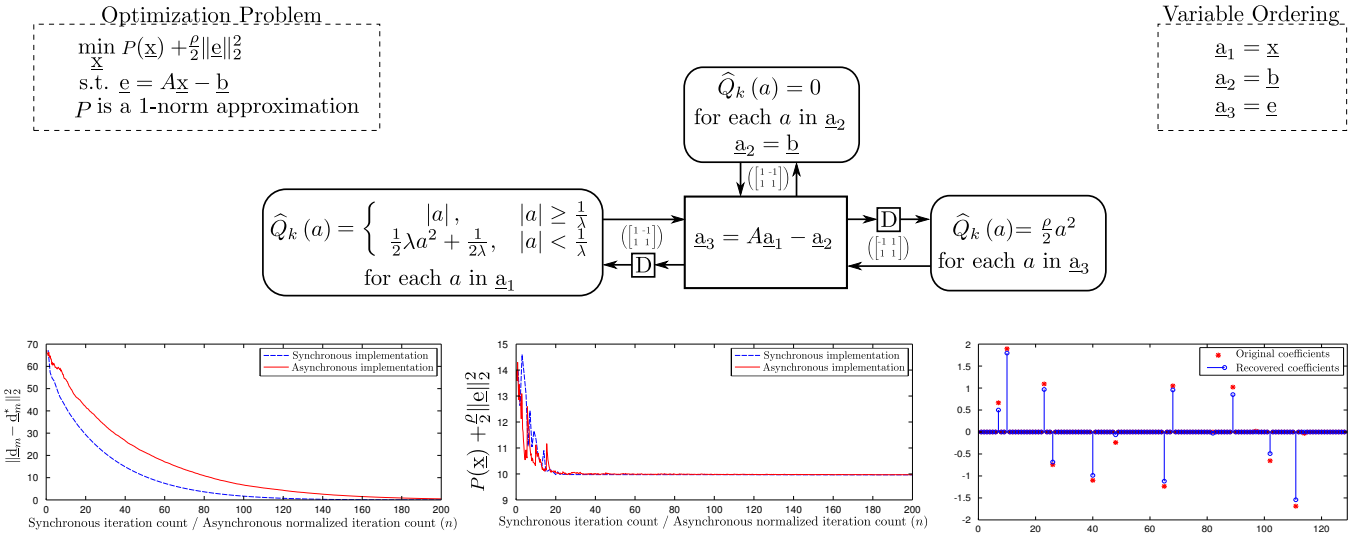
Optimization Problem
$\min_{\underline{x}} P(\underline{x}) + \frac{\rho}{2}\|\underline{e}\|_2^2$
s.t. $\underline{e} = A\underline{x} - \underline{b}$
P is a 1-norm approximation
Variable Ordering
$\underline{a}_1 = \underline{x}$
$\underline{a}_2 = \underline{b}$
$\underline{a}_3 = \underline{e}$
$\widehat{Q}_k(a) = 0$ for each a in $\underline{a}_2$
$\underline{a}_2 = \underline{b}$
$\widehat{Q}_k(a) = \begin{cases} |a|, & |a| \geq \frac{1}{\lambda} \\ \frac{1}{2}\lambda a^2 + \frac{1}{2\lambda}, & |a| < \frac{1}{\lambda} \end{cases}$ for each a in $\underline{a}_1$
$\underline{a}_3 = A\underline{a}_1 - \underline{a}_2$
$\widehat{Q}_k(a) = \frac{\rho}{2}a^2$ for each a in $\underline{a}_3$
D
Synchronous implementation
Asynchronous implementation
$\|\underline{d}_m - \underline{d}_m^\star\|_2^2$
Synchronous iteration count / Asynchronous normalized iteration count (n)
$P(\underline{x}) + \frac{\rho}{2}\|\underline{e}\|_2^2$
Original coefficients
Recovered coefficients

Fig. 2. Signal processing architecture and numerical simulation corresponding to an algorithm for solving the LASSO problem. An approximation to the 1-norm is used that is quadratic in the close vicinity of 0. The parameters λ and ρ are selected to specify the interval outside of which the 1-norm approximation is exact and to trade off between the sparsity of the solution and the enforcement of the linear equality constraints, respectively. For the depicted solution λ and ρ are selected to be large. Note in particular the monotonic convergence of $||\underline{d}_m - \underline{d}_m^\star||_2^2$ to zero. "Asynchronous normalized iteration count" indicates the number of iterations times the probability of sampling, discussed in Section II.

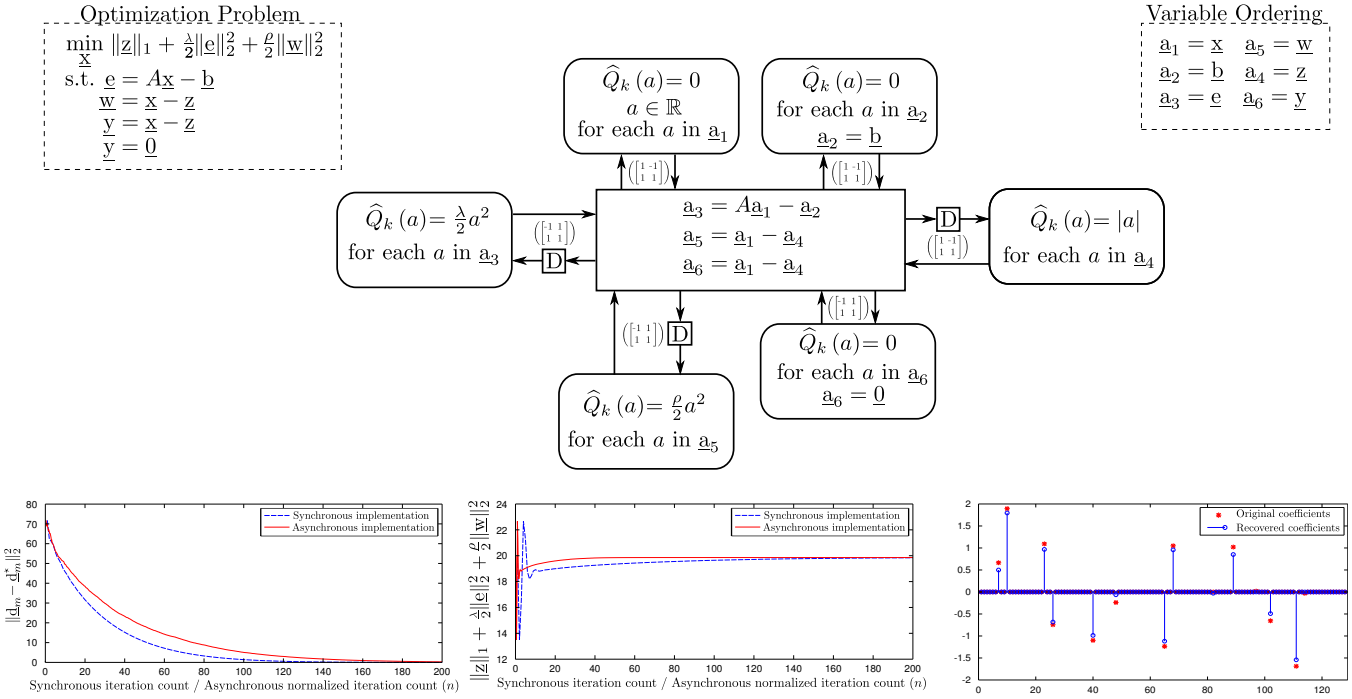
Optimization Problem
$\min_{\underline{x}} \|\underline{z}\|_1 + \frac{\lambda}{2}\|\underline{e}\|_2^2 + \frac{\rho}{2}\|\underline{w}\|_2^2$
s.t. $\underline{e} = A\underline{x} - \underline{b}$
$\underline{w} = \underline{x} - \underline{z}$
$\underline{y} = \underline{x} - \underline{z}$
$\underline{y} = \underline{0}$
Variable Ordering
$\underline{a}_1 = \underline{x}$ $\underline{a}_5 = \underline{w}$
$\underline{a}_2 = \underline{b}$ $\underline{a}_4 = \underline{z}$
$\underline{a}_3 = \underline{e}$ $\underline{a}_6 = \underline{y}$
$\widehat{Q}_k(a) = 0$ $a \in \mathbb{R}$ for each a in $\underline{a}_1$
$\widehat{Q}_k(a) = 0$ for each a in $\underline{a}_2$ $\underline{a}_2 = \underline{b}$
$\widehat{Q}_k(a) = \frac{\lambda}{2}a^2$ for each a in $\underline{a}_3$
$\underline{a}_3 = A\underline{a}_1 - \underline{a}_2$
$\underline{a}_5 = \underline{a}_1 - \underline{a}_4$
$\underline{a}_6 = \underline{a}_1 - \underline{a}_4$
$\widehat{Q}_k(a) = |a|$ for each a in $\underline{a}_4$
$\widehat{Q}_k(a) = 0$ for each a in $\underline{a}_6$ $\underline{a}_6 = \underline{0}$
$\widehat{Q}_k(a) = \frac{\rho}{2}a^2$ for each a in $\underline{a}_5$
D
Synchronous implementation
Asynchronous implementation
$\|\underline{d}_m - \underline{d}_m^\star\|_2^2$
Synchronous iteration count / Asynchronous normalized iteration count (n)
$\|\underline{z}\|_1 + \frac{\lambda}{2}\|\underline{e}\|_2^2 + \frac{\rho}{2}\|\underline{w}\|_2^2$
Original coefficients
Recovered coefficients

Fig. 3. Signal processing architecture and numerical simulation for an augmented-cost LASSO problem, with the cost being augmented similarly to that of various ADMM formulations.[3] The augmentation parameter is denoted ρ. The parameter λ is selected to trade off between the sparsity of the solution and the enforcement of the linear equality constraints. For the depicted solution λ and ρ are selected to be large.

except at 0. Eq. 2 implies, for example, that coupling the constitutive relations denoted $m(\cdot)$ to the linear interconnection elements via deterministic vector delays, the discrete-time signal denoted $\underline{d}'_m[n]$ will converge to 0 and so the signal $\underline{d}_m[n]$ will converge to $\underline{d}_m^\star$.

The uniqueness of the stationarity conditions and the property of the constitutive relations being dissipative used in the preceding argument are not, however, strictly required. A more general line of reasoning involves justifying Eq. 2 in the vicinity of any such solution $c_i^\star$ and $d_i^\star$, for example by claiming that even if specific constitutive relations $m_k(\cdot)$ are norm-increasing, the overall interconnected system results in a map from $\underline{d}'_m$ to $\underline{d}_D$ that is norm-reducing in the vicinity of that solution.

Arguments for convergence involving essentially Eq. 2 can also be applied in a straightforward way to systems utilizing asynchronous delays, modeled as discrete-time sample-and-hold systems triggered by independent Bernoulli processes. In particular by taking the expected value of $||\underline{d}'_{(m)}[n]||$, applying the law of total expectation, substituting in Eq. 2, and performing algebraic manipulations, it can be argued that $E[||\underline{d}'_{(m)}[n]||]$ converges to 0. A more formal treatment of convergence is the subject of future work.

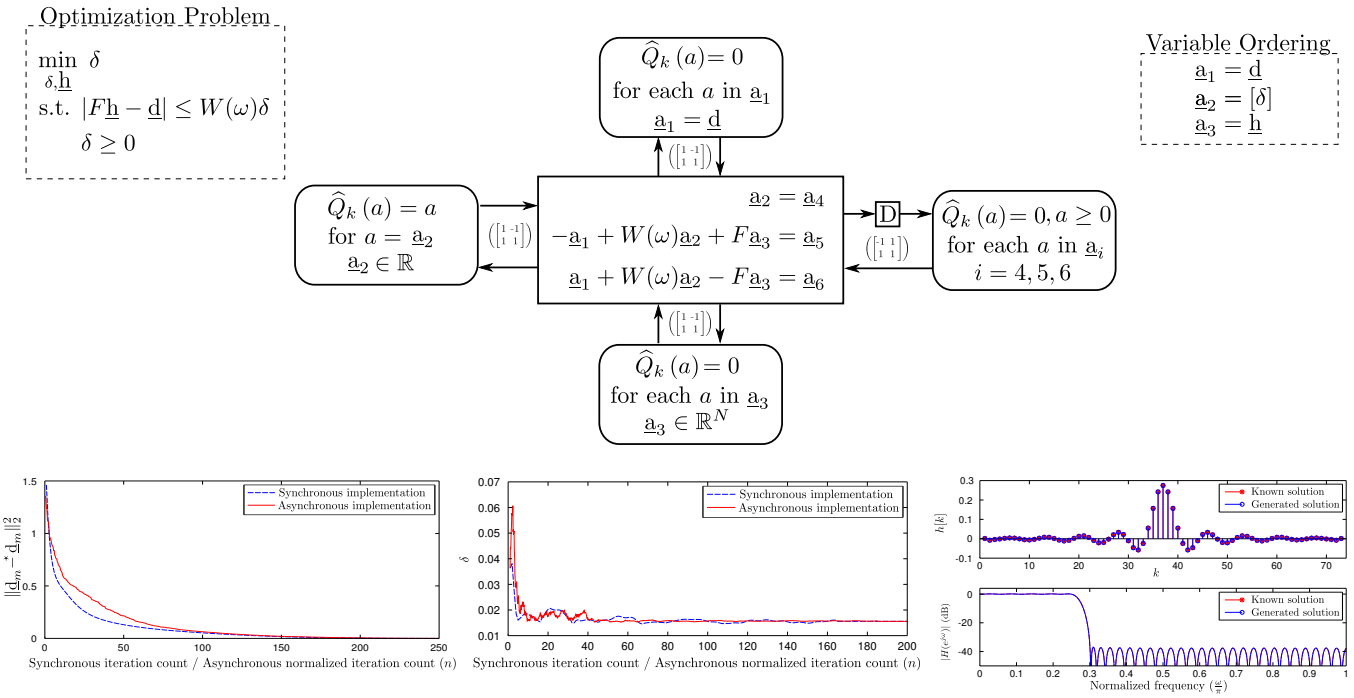

Fig. 4. Signal processing architecture and numerical simulation corresponding to a minimax-optimal FIR filter design problem, specifically that of lowpass filter design. The obtained result is compared with a known solution from the Parks-McClellan algorithm.

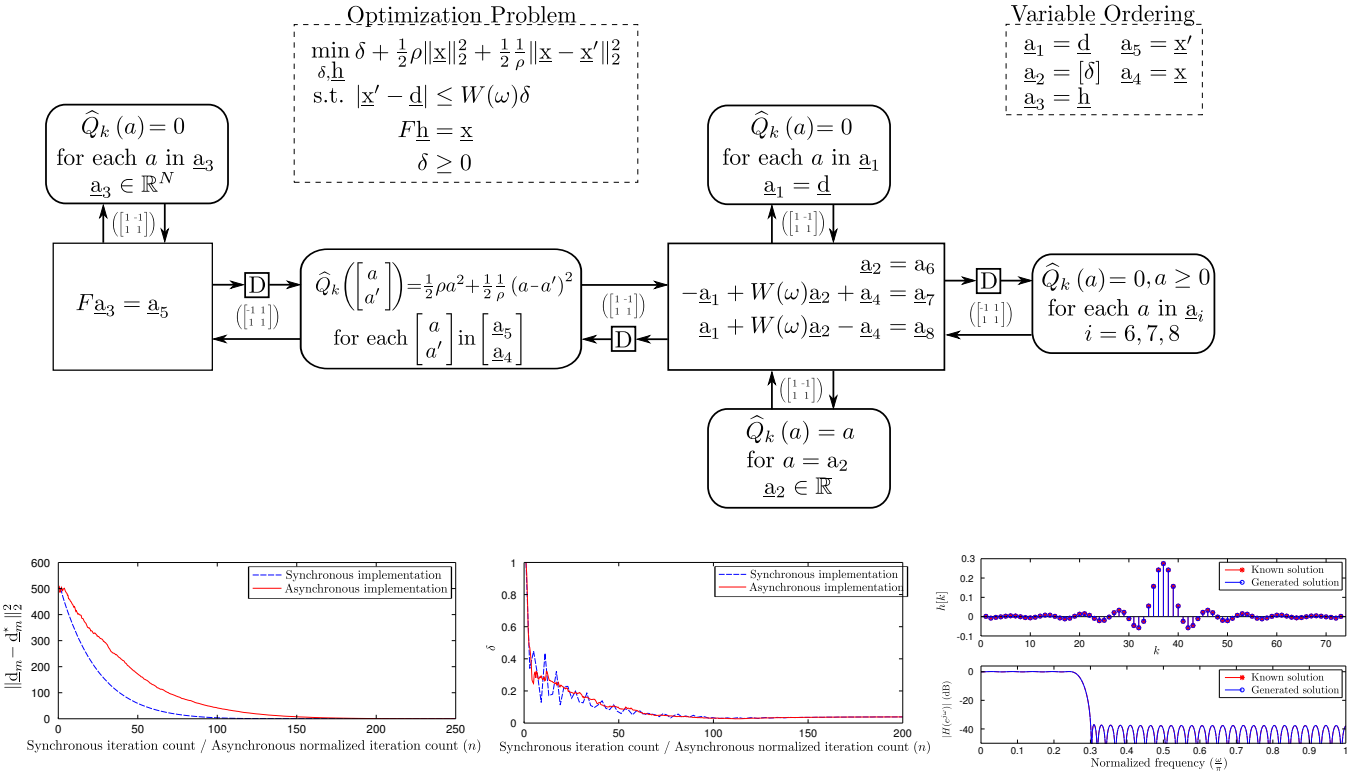

Fig. 5. Alternative algorithm for minimax-optimal filter design, obtained by modification of the problem statement in Fig. 4 and intended to demonstrate that the presented framework can be used in designing a variety of distinct classes of algorithms. The parameter ρ is selected to specify the relative enforcement of equality between the system variables loosely shared between the two linear interconnection elements. For the depicted solution ρ is selected to be small, resulting in a very close approximation to the lowpass filter design problem in Fig. 4.

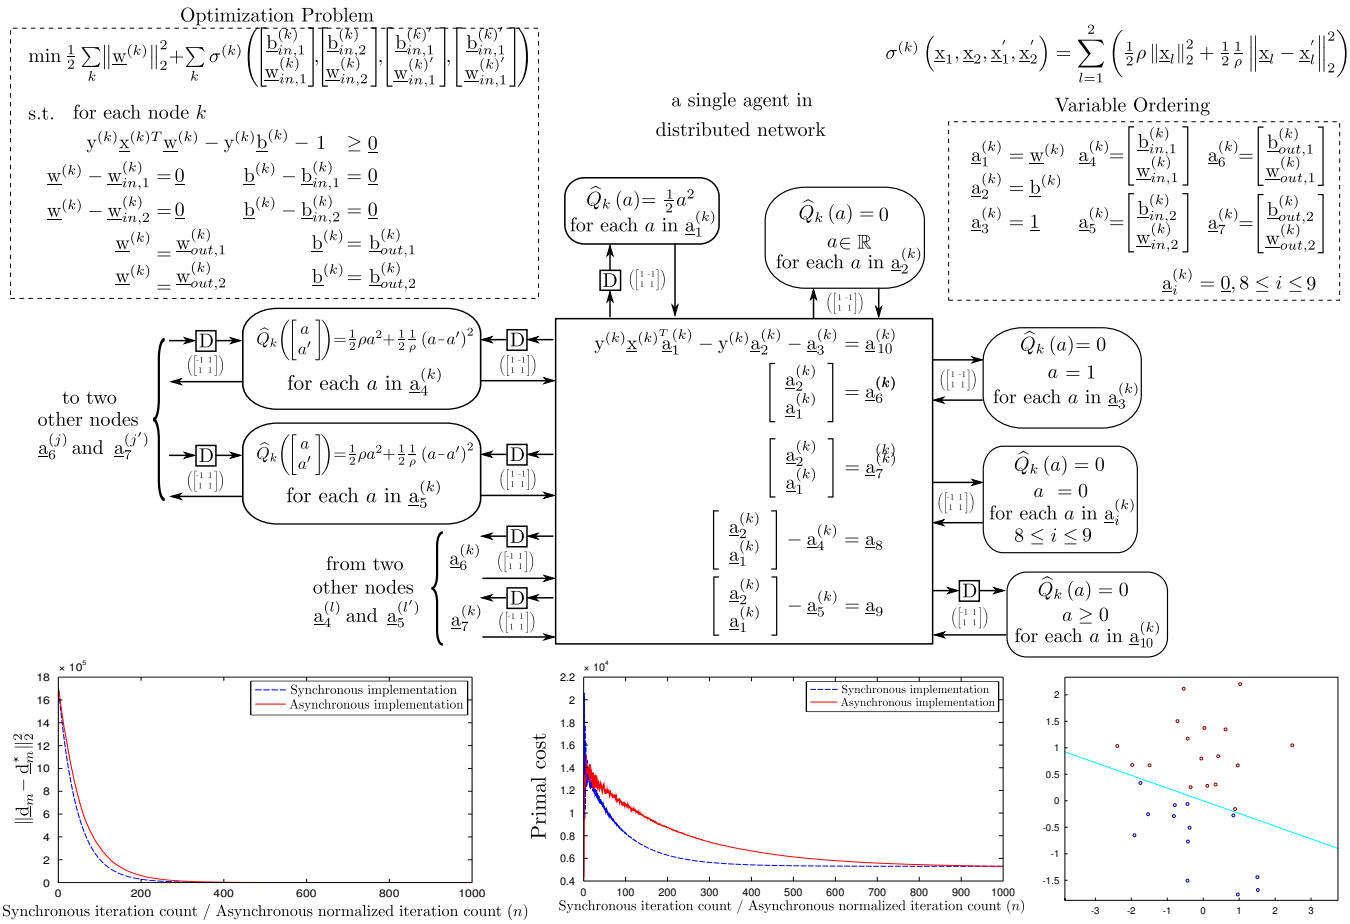

Fig. 6. Signal processing architecture for a single agent in a connected graph implementing a decentralized algorithm for training a support vector machine classifier. The numerical simulation depicts a system involving 30 such agents, each having knowledge of a single training vector. The parameter ρ specifies the relative enforcement of equality for the system variables that are coupled between each agent in the graph. For the depicted solution ρ is selected to be small, and the graph is known to be connected, with each node as depicted above having exactly four incident connections.

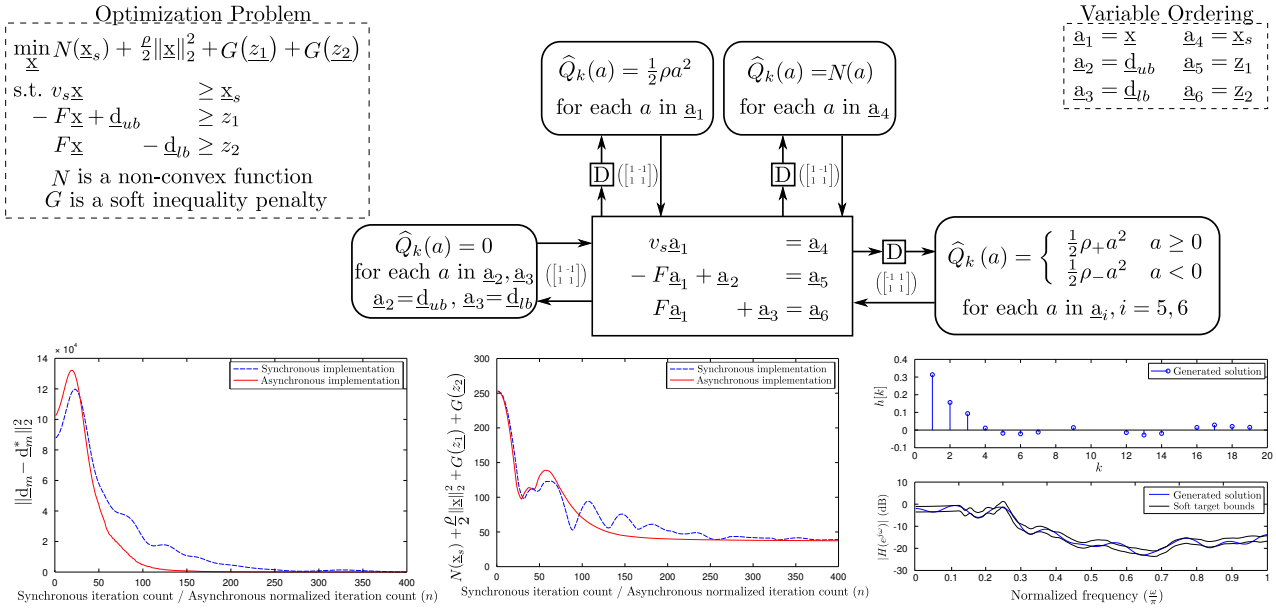

Fig. 7. Signal processing architecture and numerical simulation corresponding to a nonconvex sparse filter design problem. The parameters ρ and v_s are respectively selected to specify the enforcement of the size of $\underline{x}$ and the width of the abrupt decrease in cost about 0 for the nonconvex element. ρ_+ and ρ_- affect the enforcement of the soft inequality constraints. For the depicted solution ρ and ρ_+ are selected to be small and ρ_- and v_s are selected to be large.

REFERENCES

[1] T. A. Baran and T. A. Lahlou, "Conservative signal processing architectures for asynchronous, distributed optimization part I: General framework," in *Proc. of IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, 2014.

[2] T. Baran, D. Wei, and A. V. Oppenheim, "Linear programming algorithms for sparse filter design," *Signal Processing, IEEE Transactions on*, vol. 58, no. 3, pp. 1605–1617, 2010.

[3] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, 2011.

[4] T. A. Baran, *Conservation in Signal Processing Systems*, Ph.D. thesis, Massachusetts Institute of Technology, 2012.

[5] A. Fettweis, "Wave digital filters: Theory and practice," *Proceedings of the IEEE*, vol. 74, no. 2, pp. 270–327, Feb 1986.

[6] L. O. Chua and G. N. Lin, "Nonlinear programming without computation," *Circuits and Systems, IEEE Transactions on*, vol. 31, no. 2, pp. 182–188, Feb 1984.

[7] J. B. Dennis, *Mathematical Programming and Electrical Networks*, Ph.D. thesis, Massachusetts Institute of Technology, 1958.

[8] M. P. Kennedy and L. O. Chua, "Neural networks for nonlinear programming," *Circuits and Systems, IEEE Transactions on*, vol. 35, no. 5, pp. 554–562, May 1988.

[9] W. Millar, "Some general theorems for non-linear systems possessing resistance," *Philosophical Magazine Series 7*, vol. 42, no. 333, pp. 1150–1160, 1951.

[10] J. Wyatt, "Little-known properties of resistive grids that are useful in analog vision chip designs," *Vision Chips: Implementing Vision Algorithms with Analog VLSI Circuits*, pp. 72–89, 1995.

[11] P. Penfield, R. Spence, and S. Duinker, *Tellegen's Theorem and Electrical Networks*, The MIT Press, 1970.

[12] B. D. H. Tellegen, "A general network theorem, with applications," Tech. Rep., Philips Research Reports, Philips Research Reports.

[13] J. C. Willems, "Dissipative dynamical systems part I: General theory," *Archive for Rational Mechanics and Analysis*, vol. 45, pp. 321–351, jan 1972.

[14] J. C. Willems, "The behavioral approach to open and interconnected systems," *Control Systems, IEEE*, vol. 27, no. 6, pp. 46–99, Dec 2007.

[15] T. A. Baran and B. K. P. Horn, "A robust signal-flow architecture for cooperative vehicle density control," in *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, May 2013, pp. 2790–2794.

[16] S. K. Rao and T. Kailath, "Orthogonal digital filters for vlsi implementation," *Circuits and Systems, IEEE Transactions on*, vol. 31, no. 11, pp. 933–945, Nov 1984.

[17] E. Deprettere and P. Dewilde, "Orthogonal cascade realization of real multiport digital filters," *International Journal of Circuit Theory and Applications*, vol. 8, no. 3, pp. 245–272, 1980.

[18] N. Parikh and S. Boyd, "Block splitting for distributed optimization," *Mathematical Programming Computation*, 2014.

[19] P. A. Forero, A. Cano, and G. B. Giannakis, "Consensus-based distributed support vector machines," *J. Mach. Learn. Res.*, 2010.

[20] E. Wei and A. Ozdaglar, "Distributed alternating direction method of multipliers," in *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, Dec 2012, pp. 5445–5450.